Advances in the Dynallax Solid-State Dynamic Parallax Barrier Autostereoscopic Visualization Display System

Tom Peterka[1], Robert L. Kooima[2], Daniel J. Sandin[2,3], Andrew Johnson[2], Jason Leigh[2], Thomas A. DeFanti[2,3]

Abstract— A solid-state dynamic parallax barrier autostereoscopic display mitigates some of the restrictions present in static barrier systems, such as fixed view-distance range, slow response to head movements, and fixed stereo operating mode. By dynamically varying barrier parameters in real time, viewers may move closer to the display and move faster laterally than with a static barrier system, and the display can switch between 3D and 2D modes by disabling the barrier on a per-pixel basis. Moreover, Dynallax can output four independent eye channels when two viewers are present, and both head-tracked viewers receive an independent pair of left-eye and right-eye perspective views based on their position in 3D space. The display device is constructed by using a dual-stacked LCD monitor where a dynamic barrier is rendered on the front display and a modulated virtual environment composed of two or four channels is rendered on the rear display. Dynallax was recently demonstrated in a small-scale head-tracked prototype system. This paper summarizes the concepts presented earlier, extends the discussion of various topics, and presents recent improvements to the system.

Index Terms— I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism --- virtual reality. Keywords: autostereoscopic display, Dynallax, parallax barrier, Varrier, visualization, 3D display

1 Introduction

Lenticular screens and parallax barrier strip displays are dominant autostereoscopic (AS) technologies [7] [13]. Published literature on AS technology abounds, and any internet search quickly reveals the myriad commercial products that are available. One of the few barrier strip AS display systems suitable for production-quality scientific visualization is Varrier [11] [20], from the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago. Most users of Varrier are satisfied with the overall static barrier AS experience, but can also quickly point out some shortcomings: limited spatial resolution, restrictions on head movement speed, and the fact that only one user can view the AS display at a time.

A dynamic parallax barrier, Dynallax, addresses some of those deficiencies. Moreover, it affords other benefits, including expanded view distance working range, reduced sensitivity to system latency during head movement, eliminated physical barrier registration, ability to disable the barrier and convert the display to 2D, and the affordance of two independently tracked viewers, each with its own AS perspective of the virtual world.

[1] *Argonne National Laboratory*
[2] *Electronic Visualization Laboratory, University of Illinois at Chicago*
[3] *California Institute for Telecommunications and Information Technology University of California at San Diego*

address correspondence to tpeterka@mcs.anl.gov

Manuscript received October 9, 2007

2 Background

The function of a parallax barrier is to occlude certain regions of an image from each of the two eyes, while permitting other regions to be visible, as in Fig. 1. By simultaneously rendering strips of a left eye image into the regions visible by the left eye and likewise for the right eye, a complete perspective view is directed into each eye. Then, by fusing two stereo images into a 3D representation, an AS experience results without the need for 3D glasses.

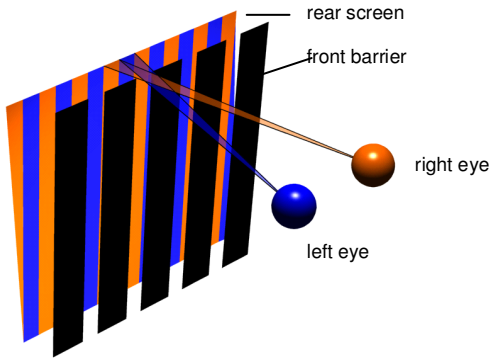

Figure 1: A parallax barrier is a series of transparent and opaque strips that permit each eye to see different regions of the display screen behind the barrier.

When the AS concept is coupled with real-time view-update, head-tracking, first-person perspective, and interactive application control, an AS virtual reality (VR) system results. Figure 2 illustrates two implementations of the Varrier (static barrier) system in large and small versions.

The parallax barrier is a high-resolution printed film that is affixed to a glass substrate and appended to the front of an LCD monitor [12]. Another popular variation is the lenticular screen. Both function equivalently [5]. The printed pattern on the barrier is a fine-pitch sequence of opaque and transparent strips; the period of this repeating pattern is on the order of .5 to 1 mm.

In a static barrier, the period of this barrier is determined a priori by the system designer, and such early-stage design decisions impact primary system responses that cannot be varied once built. System outputs such as view distance operating range (minimum, maximum, optimum), visual acuity, and the fact that the system is capable of supporting only one user at a time are three such results of barrier period choice. The consequences of these design-time decisions are magnified by the long turn-around time to correct or modify the barrier. Moreover, with respect to supporting two tracked viewers, no single optimal barrier period can be preselected; in Sections 3 and 4 it will be shown that the barrier period must vary in real time to prevent interchannel conflicts.

In addition to fixed working range and strict single-user mode for tracked two-view systems, static barrier AS systems have other disadvantages. One is that the barrier cannot be disabled, affording a convertible 3D - 2D system such as [14]. Moreover, such systems are limited by the fact that performance criteria such as frame rate and latency are more critical in fixed-barrier AS systems than in other stereo techniques. Unlike passive and active stereo VR, rapid head movements result not only in scene lag but also in visible artifacts when the head velocity is faster than channels can be updated. Defects such as image flicker, black banding, and ghosting are visible in a head-tracked AS VR system during head movements and disappear when the viewer stops moving.

Mechanical dynamic parallax barriers have appeared in the AS literature. For example, Moseley et al. [16] proposed an electro-mechanical combination of barriers to produce dynamic results in response to head-tracked viewer positions. Solid-state dynamic barriers also have appeared in the literature, but the barrier architecture differs from Dynallax and serves a different purpose. For example, the Cambridge Display [17] time-multiplexes a number of channels to predetermined, untracked positions in space using a ferroelectric liquid crystal shutter.

Figure 2: The Varrier system is shown in large tiled (left) and smaller desktop (right) versions.

The NYU display [18] utilizes a rapidly oscillating barrier rendered on a pi-cell ferroelectric liquid crystal shutter. By combining spatial multiplexing inherent in a parallax barrier with the time multiplexing of three barrier phases, this system increases spatial resolution and simultaneously conceals the coarse scale of the barrier strips.

3 System Description

3.1 Dynamic barrier construction

A dual-stacked LCD display can be constructed by placing a second liquid crystal (LC) layer in front of the first, such that both LC layers are illuminated by a single backlight, and the front layer angle of polarization is orthogonal to the rear. Equation 1 shows that the final output intensity I_{final}, of such a display, is the product of the two individual layer intensities, I_{rear} and I_{front}. In other words, the light generated by the backlight is attenuated by the rear LC layer, and then this result is further attenuated by the front LC layer.

$$I_{final} = I_{rear} \cdot I_{front} \qquad (1)$$

A complete dual-stacked display is available from PureDepth [9], and this is the display used for the Dynallax prototype.

3.2 System structure

Dynallax is structured as a set of three processes: a master controller, the front screen rendering slave, and the rear screen rendering slave. Interprocess communication is accomplished with MPICH-2 [6]. Originally each process resided on a separate Linux machine, but the current architecture is more economical and compact, consisting of a single machine with two separate graphics cards, not linked by SLI. The first card drives the front screen and an optional console, while the second card is dedicated to the rear screen. The front and rear screens of the dual-stacked display cannot be driven from the same graphics card because some Dynallax modes require asynchronous update of the two screens.

Head tracking is accomplished by using an Intersense 900 [2] tracking system with two sensors, for two viewers. Each viewer wears one sensor on a headband on the forehead. Later implementations will utilize tetherless camera tracking, as in [1]. At this point in the research, however, tethered sensors are sufficient for testing and as a proof of concept, as in Figure 3, which depicts one user wearing a tracking sensor on a headband and interactively navigating through a VR scene.

An independent first person-perspective is channeled to each viewer (up to two). Real-time interactivity is afforded with either a tracked or untracked wand containing a joystick and multiple buttons. In the case of two viewers, the wand is either shared or controlled by only

Figure 3: Wearing a tracking sensor on a headband, a user interacts with Dynallax. Eventually, tracking will be camera-based and tetherless.

one viewer at a time. Most of this paper, however, is not dedicated to describing system hardware or tracking and navigation controls. This information is included for the sake of completeness, but the main advances of this research are the algorithms that compute the scene and barrier images, and how these algorithms change the images depending on actions of the viewer or viewers.

3.3 Image Computation Algorithm

We turn our attention now to the method of computing the barrier pattern on the front screen and the interleaved images on the rear screen. For the time being, let us assume only a single viewer, so that two eye channels are spatially multiplexed on the rear screen. The extension to two viewers and four eye channels occurs in Section 3.4.3. First, we introduce some underlying ideas and motivations, before describing the solution in greater depth.

The barrier pattern is a sequence of alternating transparent and opaque strips, as in Figure 1. On the front screen, one period is made up of a single transparent-opaque cycle, whereas on the rear screen, one period contains one channel for each eye. In the following discussion, we often refer to a single barrier period rather than the entire pattern, with the understanding that conclusions for a single period transfer readily to an entire barrier pattern.

Since the viewing model in Dynallax is based on perspective projection, as opposed to parallel projection, the size and position of any barrier period on the front screen is related to the corresponding barrier period on the rear screen by a perspective transformation. In the general case, perspective projection is nonlinear, but in this special case where the two projection planes are parallel, the projection degenerates to a convenient linear transformation. This is shown in Figure 4, where the expressions for scaling and shifting the barrier period from front to rear screen are given in terms of p, the front screen barrier period, p', the rear screen barrier period, t, the optical thickness between the two screens, d, the distance from the eye plane to the front screen plane, x, the eye position in the x direction, and x', the rear screen barrier shift. To summarize, the barrier pattern on the front screen is for the purpose of visibly rendering the pattern, while the barrier pattern on the rear screen is for the purpose of composing image channels. In either case, however, a computational model of the pattern must be maintained, and the two models are related by a scaling and translation operation.

Dynallax is motivated by our previous work with static barrier autostereo barrier models, in particular the Autostereo Combiner work of Kooima et al. [3], [4]. In such a subpixel barrier model, the size of a channel is not limited by the size of a pixel, and is often smaller. We chose to replicate the main concepts of that model in Dynallax, in order to leverage the advantages of a subpixel barrier model such as increased visual acuity and net effective resolution. In discussing subpixel barrier models, it should be noted that two entirely different approaches are possible, an integer-valued (or discrete) model, or a real-valued (or continuous) model. The real-valued model is the one employed in Autostereo Combiner and in Dynallax.

In the integer-valued model, the barrier is stored in a 2D image. Such a subpixel model is proposed in earlier literature [13], [15], and can be implemented in graphics as a texture where various discrete subpixel size windows are illuminated while others are not, as in Figure 5. This texture is registered with the underlying display substructure. The integer-valued model was tested initially for Dynallax, but abandoned due to quantization errors that occur when scaling and translating the front barrier model to the rear barrier model. These aliasing artifacts are unavoidable when an already discrete model is discretized (i.e., sampled) a second time during that transformation.

Rather, a real-valued function is defined to model the barrier, defined by the following five parameters:

1. Barrier period p or line spacing in lines per unit distance
2. Angular orientation θ (line tilt) in degrees from vertical
3. Duty cycle D (fraction of opaque to total period)
4. Lateral shift s (left-right) along the plane of the display screen from some reference point
5. Optical distance t that the barrier is located in front of the rear display screen

The last two parameters can be considered as the position of the barrier in 3D space, and thus the barrier is defined by its position, orientation, period, and duty cycle, and all of these parameters are elements of the real number domain. This design offers several advantages. For example, not only is the barrier period continuously adjustable, but so is the line tilt angle; hence barrier parameters are infinitely variable to dynamically optimize viewing conditions. Moreover, when the barrier is transformed by perspective projection from the front to the rear screen, quantization errors do not result.

In fact, the graphics card driver's antialiasing features are not necessary and are typically disabled in Dynallax. Since both front and rear screens are maintained in real-valued functions, discretization occurs exactly once for

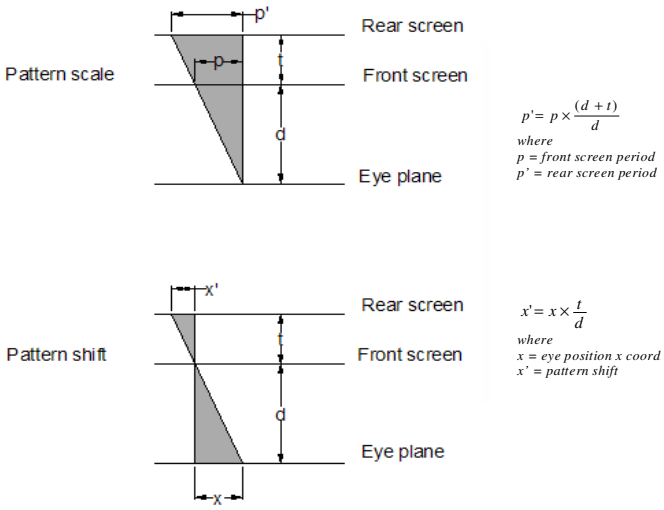

Figure 4: The barrier period and position between front and rear screens are related by a scale and shift operation.

each screen, late in the rendering pipeline, when the front and rear screen images are rasterized by the graphics hardware. A logical future research direction may be to search for alternative barrier patterns besides alternating strips. The implication of this current research is to limit that search space to only functions that can be described by real-valued parameters. Barrier models that can only be defined as discrete images are inappropriate for Dynallax.

Dynallax's barrier model is conceptually equivalent to the Autostereo Combiner algorithm [3], [4]. The main differences between the algorithms is that Dynallax always uses a constant quality level of 1.0 [4] and twice as many eye channels are supported, permitting two viewers to modulate a total of four eye channels. Both models at their lowest level consider a barrier period as a step function, where 0 corresponds to opaque and 1 corresponds to clear. The step function representation of a horizontal "slice" through the barrier appears in Figure 6, along with some of the parameters that define it, such as D, the duty cycle, p', the scaled period, and θ, the line tilt angle. The duty cycle is the fraction of a period that is opaque, usually .75. The scaled period is due to perspective projection of Figure 4, and the line tilt angle refers to the fact that barrier lines are oriented at a non-vertical angle on the screen.

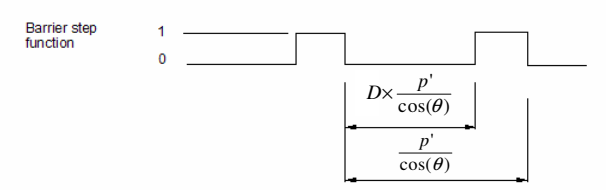

where p' = scaled barrier period, D = duty cycle (e.g. .75), θ = barrier line tilt angle

Figure 6: The barrier step function is shown.

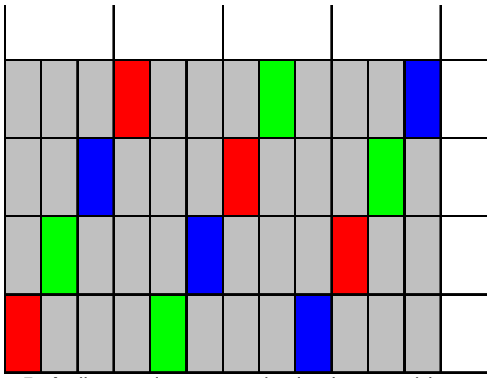

Figure 5: A discrete image can be implemented in a texture as shown above, but this is not compatible with Dynallax and results in marked quantization artifacts.

The barrier period is continuously variable down to a theoretical minimum of one subpixel. In practice however, more restricted limits are imposed on the usable barrier range by the physical hardware, not by the algorithm. At one extreme, barrier periods larger than 12 pixels (channel size of 3 pixels) become visible to the viewer and divert attention from the scene to the barrier, while at the other extreme, periods smaller than 4 pixels (channel size of 1 pixel) cause noticeable ghosting because the steering of a rear colored subpixel is quantized by the color filter of the front screen to the nearest pixel only. (See Section 4.5.)

3.4 Controller algorithm

A real-time controller, shown in Figure 7, sets the barrier period p, duty cycle D, and barrier shift s at each frame update. The controller contains three modules, as in Figure 7, the view distance control, rapid steering control, and two-viewer control. Each of these modules corresponds to one of the main features of Dynallax, and is described in the following subsections.

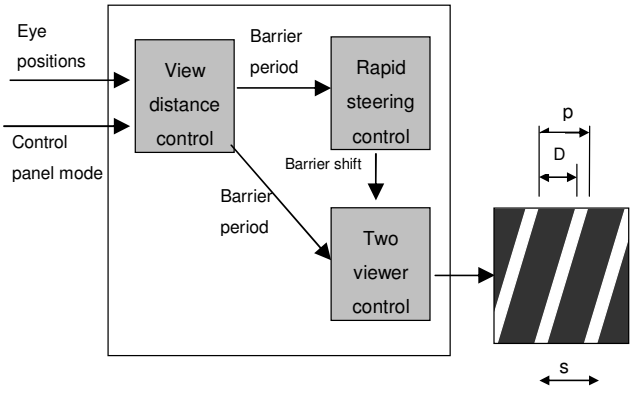

Figure 7: Controller block diagram illustrates three main functions: view distance, rapid steering, and 2-viewer control. The output of the controller are the three barrier parameters shown at the right: p = barrier period, D = duty cycle, and s = barrier shift.

3.4.1 View distance control

Optimally, we would like eye channels and guard bands (the regions separating channels) to be spaced symmetrically. In a static barrier system such as Figure 8, this regular spacing occurs at only one optimal viewing distance, determined by the barrier period.

In Dynallax, by contrast, the view distance control module sets the barrier period such that the optimal condition is always satisfied according to equation 2:

$$z_{opt} = \frac{2t(e-p)}{p} \qquad (2)$$

where
p = barrier period,
t = optical thickness between front and rear screens,
e = interocular distance, and
z_{opt} = always-optimal distance from eyes to screen.

Equation 2 shows that barrier period p is inversely proportional to view distance z_{opt}, for a given optical thickness t. This can be derived from similar triangles by casting rays from a pair of eye positions, through a pair of barrier openings, such that the rays intersect the rear screen at equidistant positions. In Dynallax, equation 2 is used to solve for p, given z, so that p is always optimal. By constantly maintaining an optimal barrier period for the current viewer distance from the screen, the system not only maintains the best possible image quality, but also permits the viewer to be closer to the screen than would otherwise be permitted by the minimum view distance for a static barrier system [12].

3.4.2 Rapid steering control

In traditional static barrier AS systems, it is easy to produce head movements that outrun the system response, resulting in momentary incorrect steering of channels to the eyes. When this situation occurs, the viewer sees black banding, because the eyes pass into the guard band regions before the scene can be updated. This is a common occurrence in Varrier, for example. The purpose of Dynallax's rapid steering feature is to weaken the dependence of stereo quality on system latency, ideally affecting only perspective correctness and not stereo quality.

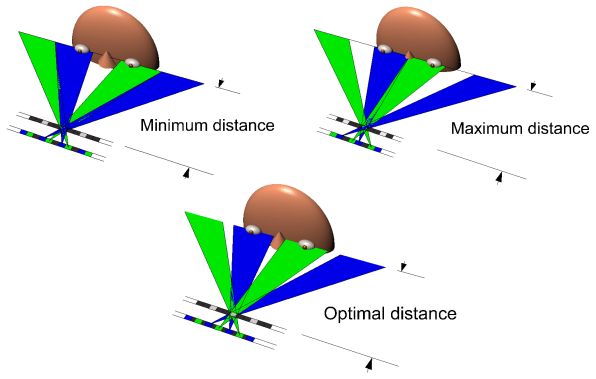

Figure 8: Spacing of channels in a static barrier system varies with view distance.

In Dynallax, the front screen has a constant computational complexity and is relatively lightly loaded, maintaining maximum frame rate, for example, 60 Hz. This is compared to the rear screen whose complexity depends on the scene complexity and drops to 15 Hz or slower for complex scenes. The reasons are social rather than technical and are a reflection of visualization usage in general. Scene complexity grows as data-set sizes increase, and frame rate is usually sacrificed for higher image quality. The key is that in Dynallax, these performance pressures impact the rear screen only, and this fact can be used to advantage.

It is exactly this disproportionate load that makes a rapid steering mode possible, where rapid head movements are accommodated by shifting the rendered front barrier rather than waiting to re-render the rear image. The rear screen continues its computation, but the front and rear screens are decoupled, allowing both screens to proceed asynchronously at their fastest rates possible. The rapid steering control module monitors this behavior and sets the barrier shift according to equation 3:

$$s = h_x \cdot \frac{t}{t+d} \qquad (3)$$

where
s = barrier shift amount,
h_x = horizontal head movement distance,
t = optical thickness between front and rear screens, and
d = normal distance from eyes to front screen plane.

Equation 3 indicates that the barrier shift is proportional to the head shift, but reduced by the ratio of distance between screens, to total distance from the head to the rear screen. This is derived by rays cast from two different head locations to the same point on the rear screen and solving the resulting similar triangle geometry.

3.4.3 Two-viewer barrier control

Problem definition

A main feature of the system is support for two tracked viewers, each receiving an independent pair of perspective eye channels. Before examining how the controller sets the barrier period in two-viewer mode, it is helpful to understand the nature of the two viewer problem.

When two viewers are present, the barrier period must increase by at least a factor of two to produce room for the additional channels to be multiplexed into the final image. Moreover, the period usually must expand further to avoid a conflict between each actual eye and the virtual repeating lobes from the other eyes. The reason is that a parallax barrier display does not only direct optical channels to the two eye positions; rather, those channels are

repeated at multiple recurring locations, or virtual lobes, across space as shown in Figure 9.

The period of repetition is given in equation 4:

$$l_x = p \cdot \frac{d + t}{t} \qquad (4)$$

where
l_x = lobe spacing in x direction,
p = barrier period,
d = normal distance from eyes to front screen plane, and
t = optical thickness between front and rear screens.

Equation 4 shows that lobe spacing l_x is directly proportional to the barrier period p, and can be derived from equation 3 by replacing barrier shift s with barrier period p, and head movement h_x with lobe spacing l_x.

General computation for n viewers

Next, we discuss how to compute the optimal barrier period in the general case of n viewers at arbitrary locations in space to prevent conflicts between eye positions and virtual lobes. A quantity called conflict energy is defined by equation 5 for each eye.

$$E_{i,j} = \begin{cases} | e - v_{i,j} |; & for\ v_{i,j} < e \\ 0; & for\ v_{i,j} >= e \end{cases} \qquad (5)$$

where
$E_{i,j}$ = conflict energy of eye$_i$ with respect to eye$_j$,
$v_{i,j}$ = distance from eye$_i$ nearest virtual lobe to eye$_j$, and
e = interocular distance.

Total conflict energy for the system is computed by equation 6, by adding the conflict energy over all pairs of relevant eye positions.

$$E_{total} = \sum_i \sum_j E_{i,j} \qquad (6)$$

Then, for any frame, the space of barrier periods is searched attempting to minimize the total system conflict energy in equation 6.

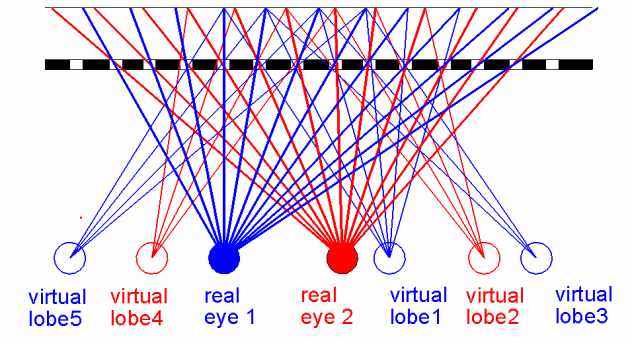

Figure 9: Repetition of channel lobes in space occurs at regularly spaced intervals. Although a limited number of lobes are shown here, this pattern continues outward in each direction.

Direct computation of two-viewer case

The previous equations were presented in [19] as a general theory for minimizing conflicts for n viewers at arbitrary locations. Fortunately, for the special case of exactly two viewers at a restricted subset of locations, a direct O(1) computation can be derived. Then, with a few minor extensions, this equation can be relaxed to accommodate two viewers at arbitrary locations, including different distances from the screen. This is the approach used in the current version of Dynallax.

The derivation begins with the most restrictive case for two viewers, that is, only one degree of freedom is permitted, x or horizontal displacement. Once this case is derived, extensions to two and three degrees of freedom will be added. We adopt the following conventions.

x = left, right location in world units
y = up, down location in world units
z = distance from front display screen, in world units
p = barrier period in world units for a single viewer
p' = barrier period in world units for two viewers
h_x = x displacement between viewers, in world units
e = interocular distance in world units
θ = barrier tilt angle

Assume that two viewers are separated by x displacement h_x such that the two-viewer barrier period $p' = 2p$. This is the optimal two-viewer condition and results in a virtual lobe separation of $2e$. As h_x increases, p' increases proportional to h_x, and screen efficiency decreases because gaps of unused space are created between lobes and channels. This pattern continues until a maximum value of $p' = 4p$, or as h_x approaches $6\ e$ in the limit from the left. At exactly $p' = 4p$, there is room for a second virtual lobe to exist between the two physical viewers, permitting p' to reset back to its optimal value of $2p$.

A pattern begins to emerge: for head displacement in the x direction only, p' follows the sawtooth non-continuous function of h_x in equation 7 and plotted in

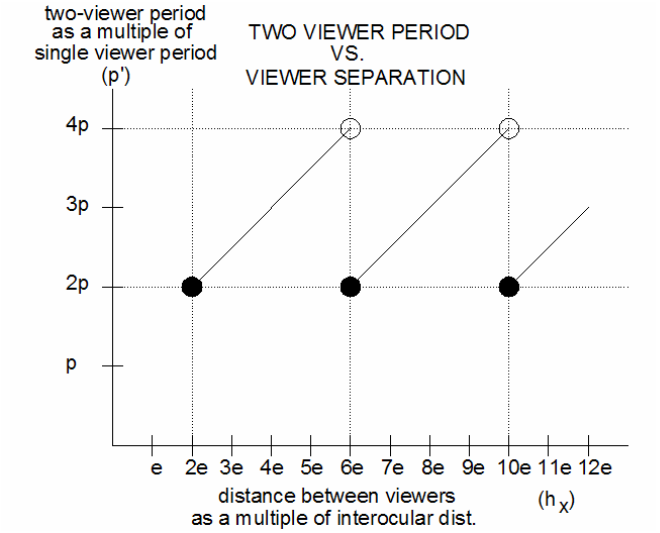

Figure 10. The barrier period function for two viewers follows a sawtooth form.

Figure 10:

$$p' = 2p + 2p \cdot fract(\frac{h_x - 2e}{4e}) \qquad (7)$$

where *fract(x)* ε (0.0, 1.0] and is the fractional part of a real-valued x. Equation 7 is the algebraic equivalent of the function in Figure 10, and the parameters of the modulo arithmetic produce the correct wavelength, amplitude, and phase shift to match that sawtooth form.

Next, the derivation is extended to permit two degrees of freedom, specifically nonequal y positions of the two viewers. The idea is to convert the actual x separation of the viewers (at different y locations) to a theoretical x separation as if the two viewers were at the same y location. One viewer is translated to the y position of the other viewer along an angle equal to the barrier tilt angle θ. The new x position is used to compute an effective h_x from equation 8,

$$h_x = \begin{cases} |\Delta eye_x| + \Delta eye_y \cdot \tan\theta & \text{for } \Delta eye_y \cdot \tan\theta > 0 \\ |\Delta eye_x| & \text{for } \Delta eye_y \cdot \tan\theta < 0 \end{cases} \qquad (8)$$

and p' is then computed from equation 7 as before. The first case of equation 8 expands the effective distance h_x according to the tangent of the line tilt angle, when that change to the effective distance is positive. In the second case, no change is performed, rather than reducing the effective distance. This is a safe behavior, so that the resulting period p' in equation 7 is never reduced, only expanded due to unequal y positions.

The last step is to permit all three degrees of freedom of viewer movement, allowing nonequal z locations. As in the previous extension, the strategy is to compute an addition to the viewer separation, Δh_x, as if the viewers were at the same z distance, and then to reuse the previous derivations from that point onward. The additional distance due to unequal z locations of the viewers is denoted by Δh_x and approximated by equation 9.

$$\Delta h_x \approx k \cdot h_x \cdot \frac{\Delta z}{z_{avg}} \qquad (9)$$

This correction adjusts the effective head separation h_x by an amount proportional to the original separation and the ratio of z difference Δz to the actual average z distance z_{avg}. The constant k is determined empirically. The quantity Δh_x forces the barrier p' to become even larger when viewing distances are unequal. Thus, equations 8 and 10 are the resulting formula for computing p' for two-viewer locations in three degrees of freedom.

$$p' = 2p + 2p\left[fract(\frac{h_x - 2e}{4e}) + k \cdot h_x \cdot \frac{\Delta z}{z_{avg}} \right] \qquad (10)$$

3.4.4 Dual-period barrier

As the barrier expands for two viewers, the system's effective resolution decreases because the barrier causes larger regions of the screen to remain black. It should be clear from the previous section that this is necessary in order to prevent conflict between virtual lobes and viewers, but a large barrier period is problematic. Large barrier lines are visible to the viewer, occlude a majority of the scene, and diminish overall brightness and resolution. Fortunately, one more improvement to the algorithm is possible that increases screen efficiency by 50% in some cases, depending on viewer locations. The idea is to change the barrier pattern from a simple repeating period to a dual period as shown in Figure 11. This pattern applies only to two-viewer mode, and only during the second half of any of the sawtooth waves in Figure 10, when the barrier period p' is greater than or equal to three times the single-viewer barrier period p.

To quantify the improvement in screen usage due to the dual-period barrier, we define the screen efficiency as the fraction of rear-screen pixels that are used to compose viewable channels. Dynallax, like Varrier, relies on guard bands to separate channels [12]. In the most efficient, optimal conditions, guard bands and channels are equidistantly arranged such that the maximum possible efficiency is .5; in other words, half of the screen is used for guard bands. Table 1 compares the efficiency of the single-period barrier and the dual-period barrier at various barrier periods under two-viewer mode. The efficiency gain in the far right column is computed as the ratio of the single-period and dual-period barrier efficiency, and represents the improvement in using a dual-period barrier over a single-period barrier.

Now that the component parts of the controller algorithm have been introduced, namely view distance, rapid steering, and two-viewer control, let us see how the result of these components affects how the master process communicates with and synchronizes the front and rear screens.

3.4.5 Interprocess communication

In order for the view distance, rapid steering, and two-viewer modules to produce their intended results, Dynallax uses both synchronous and asynchronous communication and display modes. For example, when in rapid steering mode, communication occurs asynchronously so

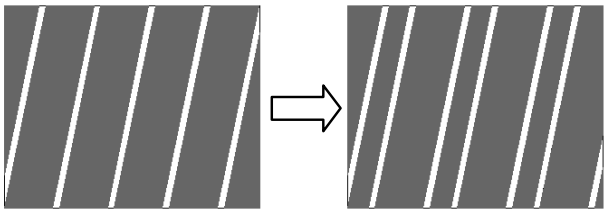

Figure 11. Changing the barrier pattern from a single period to a dual period increases efficient use of screen pixels.

TABLE 1
EFFICIENCY GAIN OF DUAL PERIOD BARRIER OVER SINGLE PERIOD BARRIER

p'	Single period efficiency	Dual period efficiency	Efficiency gain
2p	0.5	NA	NA
2.5p	0.42	NA	NA
3p	0.33	0.5	1.5
3.5p	0.29	0.44	1.55
3.999p	0.25	0.4	1.6
4p	0.5	NA	NA

that the front and rear screens update at their own maximum possible rates. However, rapid steering is disabled with two viewers present and also momentarily when the barrier period is modified by the view distance module. The front and rear screens are synchronized under these circumstances so that barrier period changes occur as seamlessly as possible.

When Dynallax determines that a barrier period change is required, the slave processes rendezvous; the change is made via synchronous communication, and then the processes continue as they were. Whenever Dynallax switches to synchronous mode, an MPICH-2 communication barrier is also employed as a secondary method of synchronizing the two slaves with each other. Figure 12 summarizes the two communication modes. The functionality in Figure 12 is executed prior to every rendered frame. Hence, the encompassing flow structure is a repeating loop, where the functions in Figure 12 constitute the body inside of that loop.

Communication in Dynallax follows a duplex handshaking protocol: slaves send ready messages to the master and the master responds by sending current data to the slaves. This flow control protocol permits slaves to run asynchronously when desired and to receive the most up-to-date information whenever they ask for it.

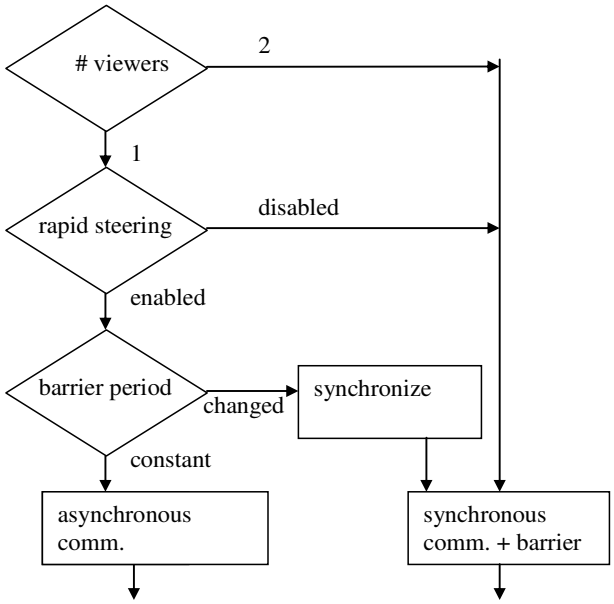

Figure 12: Interprocess communication in Dynallax is both synchronous and asynchronous.

4. RESULTS

4.1 2D, 3D, mixed mode

Dynallax is easily converted from 3D to 2D mode by rendering the front screen white, making it transparent, and rendering a single monoscopic view on the rear screen. This versatility permits the user to multiplex VR and non-VR tasks within the same display: the display space can be partitioned on a per-pixel basis to permit both tasks to be visible simultaneously. Since no glasses are required to experience the 3D scene, the user can seamlessly process both data types. For example, Figure 13 depicts mono, autostereo, and mixed modes.

4.2 View distance control

The amount of light reaching the unintended eye may be measured with a photometer and converted to a percentage of crosstalk. In a static barrier system, the crosstalk or ghost level is at a minimum at the optimal view distance and increases when the viewer approaches the near and far view distance limits of the system. This is a particular disadvantage at the near viewing limit, because viewers often prefer to be close to the screen in order to see details.

In Figure 14, the ghost level vs. view distance in the neighborhood of the near limit is plotted for two implementations of Varrier, the 35-panel Cylindrical Varrier [12] and the single-panel Personal Varrier [8]. Because the Personal Varrier system is a seated desktop display, its barrier is tuned to a smaller minimum distance than for the Cylindrical Varrier, and hence its graph is shifted to the left, but the basic pattern is the same. On this is superimposed a graph of the ghost levels for Dynallax. All three systems have a comparable minimum ghost level of

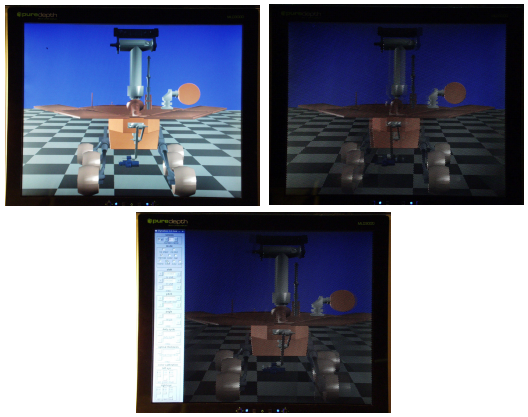

Figure 13. Dynallax permits monoscopic mode (upper left), autostereo mode (upper right), and mixed mode (2D control panel and 3D scene below).

5-7%, but the Dynallax graph does not spike upward at near distances, as do the fixed barrier systems.

It is worth investigating whether Dynallax has a far distance limit. The largest distance shown, 61 cm, is comparatively close-up for a desktop display, and Figure 14 shows that the ghost level gradually increases with distance. In theory, there should be no far limit, since the barrier period is computed optimally for any distance. However, since the period (line spacing) decreases with increasing view distance, there is a physical limitation on how narrow view channels can become based on the dot pitch and resolution of the display. Section 4.5 discusses why this limit is currently the size of a pixel when the algorithm scales down to the subpixel level. The gradual increase in ghost level is due to the nearing of this hardware display limitation, and in the current implementation, this limit occurs near the 61 cm distance.

4.3 Rapid view steering

When rapid view steering is disabled, the front and rear screens operate in lockstep with each other at the same frame rate. This is the case for two-viewer mode. On the other hand, with rapid view steering enabled, each screen is allowed to run as fast as possible and is updated with current data whenever the next frame is about to begin. The frame rate performance numbers in Table 2 bear this out for several model sizes. Images of the three test models appear below the table.

The purpose of the rapid steering mode is to relax the

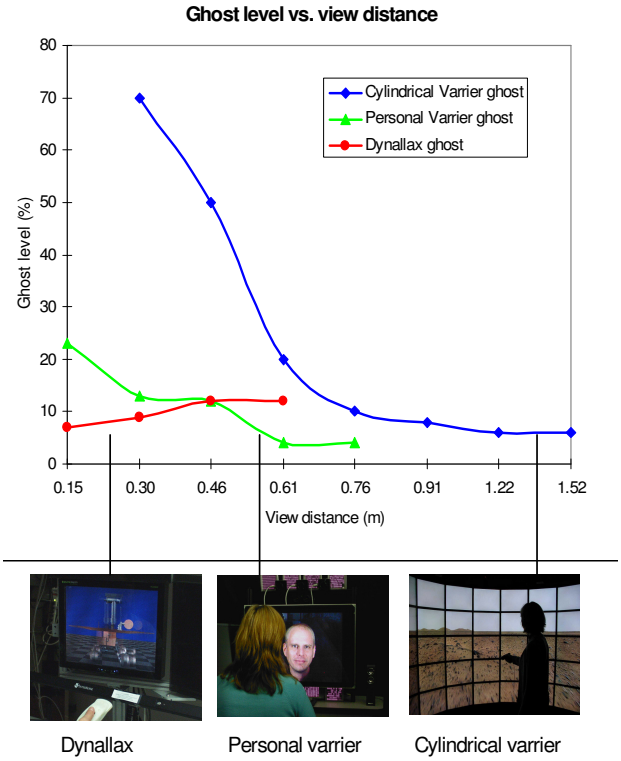

Figure 14: Ghost level vs. view distance is plotted for Dynallax and two Varrier implementations.

TABLE 2

FRAME RATES FOR FRONT AND REAR SCREENS FOR VARIOUS MODEL SIZES AND RAPID STEERING ENABLED / DISABLED

Model	# vertices	Front sync frame rate	Front async frame rate	Rear frame rate
single polygon	4	30	50	30
mars rover	15K	30	50	30
head	130K	10	50	10
skull	220K	3	50	3

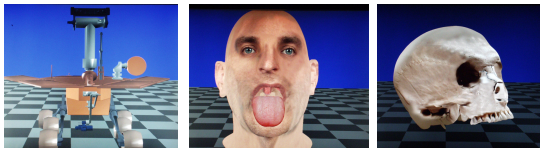

sensitivity of the parallax barrier AS system to overall system latency. For example, let us compute the maximum velocity of head movement permissible with and without the rapid steering mode, based on the skull model of Table 2. Consider a reference distance of 32 mm (half of the interocular distance) during the time that a frame is displayed. The frame time is the reciprocal of the frame rate in Table 1 plus tracker and communication latency, measured to be a constant 65 ms for a similar tracking system and cluster arrangement [12]. The last row of Table 2 indicates the following.

rapid steering: v = 32 mm / (65 + 20 ms) = .38 m / s
no rapid steering: v = 32 mm / (65 + 333 ms) = .08 m / s

In this example, the resulting speedup is a factor of greater than 4 times with rapid steering enabled.

4.4 Two viewer mode

In order to observe quality in two viewer mode, an experimental apparatus was constructed consisting of two video cameras, separated by an interocular distance of 63 mm and fixed to a jig, thereby emulating the eyes of one viewer at a known 3D position in space. In Figure 15, a test pattern of vertical, horizontal, and angled bars is rendered for the various eye channels. In other words, each different angular orientation corresponds to a different eye's view of a VR scene in normal mode. The left column of Figure 15 is the left camera (eye) view and likewise for the right column and camera (eye). In Figures 15a and 15b, only single-viewer mode is enabled, and horizontal bars are seen by the left eye and vertical bars by the right. The dim bars of the opposite orientation correspond to a ghost level of approximately 7%. This compares to Figures 15c through 15f, where two-viewer mode is enabled. Figures 15c and 15d correspond to the first viewer, while Figures 15e and 15f correspond to the second. Comparing single-viewer and two-viewer modes shows that the

ghost level decreases slightly; however, the barrier strips become more visible as their period is increased.

4.5 Other display architectures

In the current architecture purchased from PureDepth, the colored r,g,b sub-structure of the sub-pixels that make up the front screen limit the efficacy of the channel modulation algorithm when the barrier period is less than one pixel, as in Figure 16.

If the front subpixels were homogenously monochrome, for example, if the LC color filter could be removed, then the display could actually render barrier periods as low as one subpixel in channel width. Unfortunately, the color filter is an integral part of a modern LC panel. As a result, the sub-pixel resolution of the rear display is quantized to pixel resolution by the front display, since light of a given color can pass through only the same color front sub-pixel. The color filter is an integral component of a modern LC panel and cannot be removed after manufacture. However, it could be omitted during the manufacturing process.

As a transitional research step toward this architecture, the Dynallax method was tested with a monochrome

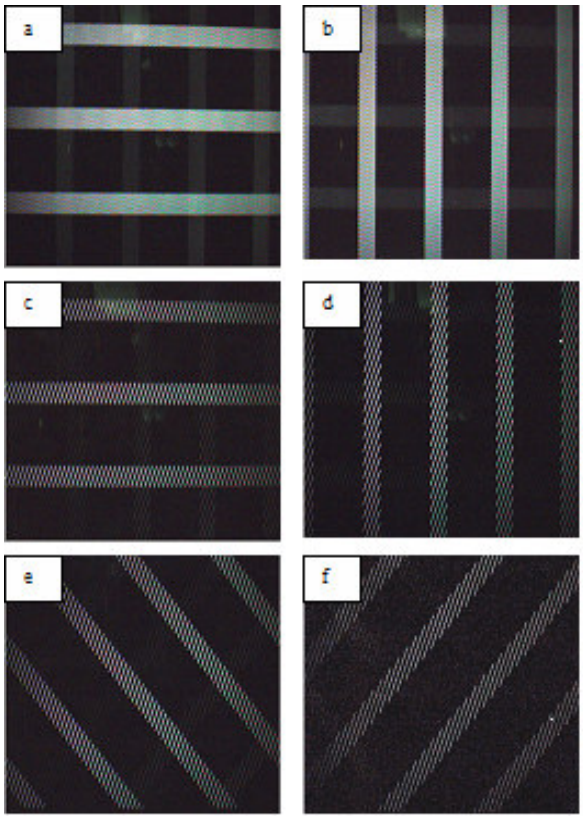

Figure 15: In this test the first viewer sees vertical and horizontal bars in the left and right eye, respectively, while the second viewer sees angled bars along two different diagonals. A pair of stereo cameras captures what either the first or second viewer would see, under single and two-viewer modes. a and b: left and right eye of single viewer mode; c and d: left and right eye of first viewer under two-viewer mode; e and f: left and right eye of second viewer under two-viewer mode. Different orientation white bars are rendered for the various channels, and the dim traces of opposite orientations are the ghost level of the system.

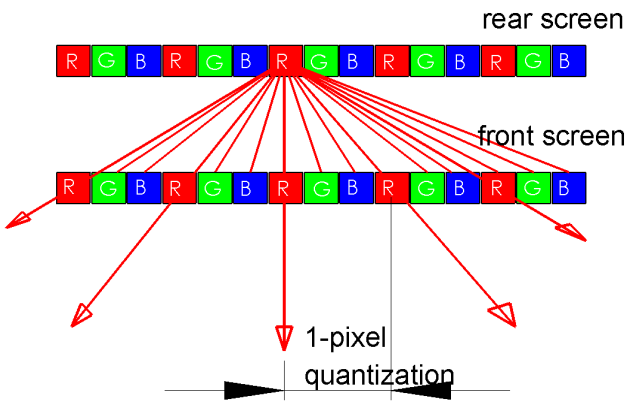

Figure 16: The colored substructure of the front display screen limits rear subpixel resolution to pixel resolution.

medical-quality LCD panel serving as the front screen of a custom-built stacked display (Figure 17). This satisfies the requirement that the front screen be monochrome; however, it is composed of pixels, not subpixels, so its dot pitch is only approximately 15% higher than the PureDepth. Still, a number of significant findings strengthen the case for the monochrome subpixel front screen in the future.

In order to physically accommodate a stacked configuration, the units are removed from their plastic and sheet metal enclosures, and various circuit boards are relocated. The angles of polarization of the two layers of a dual stacked display must be orthogonal. Rather than attempting to remove and reorient the polarizing element, the most direct approach is to simply position the entire top panel 90 degrees with respect to the lower. The top row of Figure 17 shows photographs of the custom display.

The bottom row of Figure 17 shows side-by-side comparisons of the custom display and the PureDepth display in two-viewer mode. Here, four channels are shown in different colors, photographed from far away so that all four channels are visible simultaneously. Visual quality is comparable between the two. In the custom display (lower left), the image appears only in the region where the front and rear screens overlap, since the two screens are mounted orthogonal to each other. The two screens are separated by approximately 15% smaller spacing than in the stock display (lower right), corresponding to the finer dot pitch of the front screen, resulting in proportionally smaller channel width.

This is the first time that a Dynallax system was shown to work correctly with heterogeneous front and rear display screens. The experiment shows that nothing in the Dynallax method requires the same screen size, dot pitch, or screen type. The fact that same quality results can be produced side by side with the custom and stock displays, with decreased channel size proportional to the smaller dot pitch, tells us that results should continue to scale down to monochrome subpixel dot pitch, given appropriate display hardware.

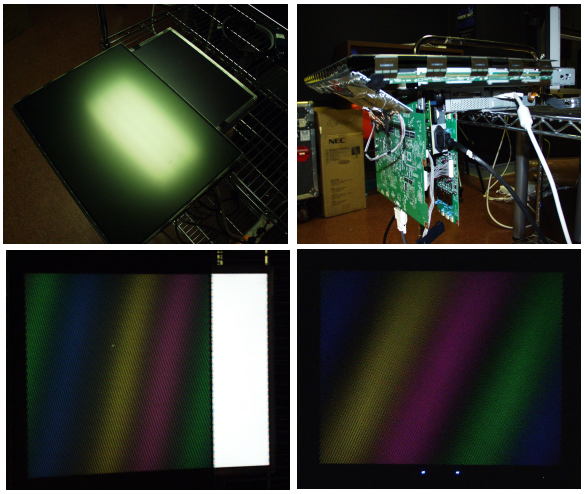

Figure 17. Photographs of the custom display are shown in the top row. The upper left photo is a top view of the second LC display panel placed on top of the first. The upper right is a view of the undersize, showing re-arranged circuitry. Results with the custom display in the lower left and the PureDepth display in the lower right are comparable and encouraging.

4.6 Net effective resolution

Net effective resolution is a parameter used to describe the output of parallax barrier AS systems. This can sometimes be a misleading metric meaning anything from a simple pixel count to a physiological measure of visual acuity, but here it is intended to count the number of minimum resolvable units (MRUs) that the system can produce. This is consistent with the approach in [20], where the MRU is equal to one barrier period horizontally by one screen pixel vertically, so that the net effective resolution is the product of the number of barrier lines and vertical pixels. In a subpixel barrier algorithm, the MRU is often some fraction of a pixel, and the purpose of a subpixel algorithm is to decrease the period size and increase the number of MRUs, thereby increasing the net effective resolution. However, Dynallax cannot arbitrarily reduce the barrier period since it is governed by the controller module that needs to satisfy constraints such as optimality and two-viewer conflict avoidance.

Table 3 demonstrates the effect of barrier size on net effective resolution under various Dynallax conditions, including single-period and dual-period two viewer conditions. These values are reported as "per eye," although it is not valid to multiply by the number of eyes in an attempt to double or quadruple the net effective resolution.

In order to compare Dynallax to other static barrier systems, [20] summarizes a number of Varrier systems in Figures 4-8 (in [20]) and in Table 1 (again in [20]). By comparison, the net resolution of Dynallax is relatively low. Although Dynallax is currently a single-panel prototype and the resolution is scalable by adding panels, Table 3 indicates the need for higher-resolution stacked displays to be produced in the future.

The base resolution of the display is 1280 x 1024, or approximately 1.3 megapixels. Table 3 shows that the most efficient use of this display, solely in terms of resolution, is in the single-viewer mode at a 610 mm view distance, or approximately 20% efficiency. By comparison, the newest generation of Varrier displays operate at 44% efficiency, computed as the ratio of net effective resolution to base display resolution. This fact reinforces the argument for a monochrome subpixel front screen that permits subpixel-size channels, since this would at least double efficiency.

5. Conclusions

Through the use of a dynamic parallax barrier, Dynallax improves the sensitivity to rapid head movement, permits two pairs of independent tracked perspective views, extends view distance range, and switches between 2D and 3D modes. In static barrier strip implementations a registration process is required to align the physical barrier, either with the underlying pixel structure or with the software's representation of the barrier computational model. This is a nontrivial time-consuming process, and the final quality of the system critically depends on the accuracy of this calibration. This task is lessened in Dynallax because it is purely a solid state, digital system.

Dynallax still has some limitations that restrict its use to research at this time. For example, the dual stacked LCD display is noticeably darker than a static barrier system; in fact both the brightness and contrast are approximately 50% that of a static system. Other drawbacks are low effective resolution and barrier line visibility. These limitations require future research before the true potential of this system can be realized. Higher resolution display devices and a monochrome subpixel front screen are necessary to reduce channel size down to the subpixel

TABLE 3
NET EFFECTIVE RESOLUTION UNDER VARIOUS CONDITIONS

Condition	View distance	Net effective resolution
1 viewer	305 mm	.13 Mpixel
1 viewer	610 mm	.26 Mpixel
2 viewers single-period optimal separation	610 mm	.13 Mpixel
2 viewers single-period worst-case separation	610 mm	.09 Mpixel
2 viewers dual-period optimal separation	610 mm	.13 Mpixel
2 viewers dual-period worst-case separation	610 mm	.07 Mpixel

scale, and this scale will increase the net effective resolution proportionally.

One method for concealing the barrier visibility is to oscillate the barrier pattern rapidly between at least two positions. This would require LC technology with a faster response time, although the trend is toward increasing refresh rate in LCD display monitors. A number of manufacturers, such as Hitachi and NEC are now introducing flat panel LCD monitors with an 8-millisecond response time [21]. The physical LC decay time has not changed; rather, the LCs are driven by a waveform of double the frequency. In order to accomplish this, a blanking interval is introduced into each cycle, and each high-voltage portion of the cycle is shorter than before. It is not clear yet whether this higher refresh frequency will support 120 Hz at full resolution; currently these monitors still support input bandwidth only up to 85 Hz.

On a related front, the ferroelectric pi-cell is a natively rapid switching LC element that has been used for many years in LC shutter glasses because of its fast response time and binary operation. Until now, however, pi-cells have been limited to monochrome micro-stripe elements rather than a matrix of fully addressable color pixels. Currently, 720 Hz ferroelectric liquid-crystal-on-silicon (FLCOS) elements are beginning to appear in color matrix form, although they are limited to small sizes [22]. These two trends, the use of rapid LCD switching cycles and the expansion of inherently fast-switching FLCOS, lead one to believe that in the future flat panel monitors may support frame rates of 120 Hz.

Dynallax is a spatial-only multiplexed method, but future fast switching display technology would permit a hybrid time-space multiplexing method. The left-right position of front barrier (and corresponding rear scene) could be rapidly alternated between at least two locations at a minimum of 120 Hz, in order to be invisible to the eye. This concept is similar to the NYU display [18][23], where a DLP projector engine in combination with a pi-cell barrier accomplished a similar result.

A second problem concerns a visible artifact, or "pop," when the barrier pattern changes scale. This is caused by imperfect synchronization between front and rear screens. The screens are currently soft-synchronized by MPICH-2, but this is not as accurate as a hardware sync signal and is subject to network performance and other factors out of our control. Unfortunately, a hard sync connection between screens is not possible because the system switches often between asynchronous and synchronous modes. Two possible solutions can be pursued.

One option is to replace the MPICH-2 communication with lower-level custom socket programming, in an effort to reduce the weight of the message passing and to take tighter control over those functions. The drawback is reduced flexibility and the need to reinvent a number of functions that already exist. The other option is to experiment with different barrier patterns besides the current set of parallel lines, in an attempt to visually "soften" the pattern. Grayscale rather than binary intensity and a random pattern of dots rather than lines are two possibilities. The idea is to conceal the existence of synchronization errors that we cannot eliminate. This is largely an unexplored area, to experiment with new patterns that have the desired characteristics of modulating eye channels while reducing their own self-visibility, and we intend to pursue this line of research in the future. (Recall from Section 3.3 that the choice of barrier pattern is limited by the fact that it must be defined by a set of real valued parameters, as opposed to a discrete pixel-based image.)

We have demonstrated a number of Dynallax's features, focusing on the system more than on applications, but we believe that the greatest advantages are yet to be realized. We foresee a wall or entire room tiled with Dynallax screens, multiple modes simultaneously active without regard to physical tile borders. These applications can include 2D text documents and powerpoints, 3D monoscopic mode scenes, 3D autostereo single viewer mode scenes, two users interacting with their own perspective of the same scene, two users viewing entirely different scenes in the same screen region (in mono or autostereo), and untracked multiview panoramagrams. At smaller form factors, we foresee pervasive Dynallax commodity desktop and laptop displays that can support 2D and 3D modes seamlessly, across applications and even mixed-mode regions within the same application. With continued research, we expect that this technology can become the ubiquitous visualization display system of the future.

ACKNOWLEDGMENTS

This work was supported in part by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357.

The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago specializes in the design and development of high-resolution visualization and virtual-reality display systems, collaboration software for use on multi-gigabit networks, and advanced networking infrastructure. These projects are made possible by major funding from the National Science Foundation (NSF), awards CNS-0115809, CNS-0224306, CNS-0420477, SCI-9980480, SCI-0229642, SCI-9730202, SCI-0123399, ANI 0129527 and EAR-0218918, as well as the NSF Information Technology Research (ITR) cooperative agreement (SCI-0225642) to the University of California San Diego (UCSD) for "The OptIPuter" and the NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement (SCI 9619019) to the National Computational Science Alliance. EVL also receives funding from the State of Illinois, General Motors Research, the Office of Naval Research on behalf of the Technology Research, Education, and Commercialization Center (TRECC), and Pacific Interface Inc. on behalf of NTT Optical Network Systems Laboratory in Japan. Varrier and CAVELib are trademarks of the Board of Trustees of the University of Illinois.

REFERENCES

[1] Javier Girado, Daniel Sandin, Thomas DeFanti, and Laura Wolf. Real-Time Camera-based Face Detection Using a Modified LAM-STAR Neural Network System. In Proceedings of SPIE-IS&T Electronic Imaging, 20-24, 2003.

[2] InterSense. InterSense IS-900 Precision Motion Tracker. http://www.isense.com/products/prec/is900/, 2006.

[3] Robert Kooima. Varrier Combiner. http://www.evl.uic.edu/rlk/varrier_combiner/varrier_combiner.html, 2006.

[4] Robert Kooima, Tom Peterka, Javier Girado, Jinghua Ge, Daniel Sandin, and Thomas DeFanti. A GPU Subpixel Algorithm for Autostereoscopic Virtual Reality. In Proceedings of IEEE Virtual Reality 2007, p. 131-138, 2007.

[5] Lenny Lipton and Mark Feldman. A New Autostereoscopic Display Technology: The SynthaGram. In Proceedings of SPIE Photonics West 2002: Electronic Imaging, San Jose, California., 2002.

[6] MPICH2. http://www-unix.mcs.anl.gov/mpi/mpich/, 2006.

[7] Takanori Okoshi. Three Dimensional Imaging Techniques. Academic Press, N.Y., 1976.

[8] Tom Peterka, Daniel Sandin, Jinghua Ge, Javier Girado, Robert Kooima, Jason Leigh, Andrew Johnson, Marcus Thiebaux, and Thomas DeFanti. Personal Varrier: Autostereoscopic Virtual Reality for Distributed Scientific Visualization. In Future Generation Computing Systems, 22, no. 8, p. 976-983, 2006.

[9] PureDepth. http://www.puredepth.com, 2006.

[10] Randi Rost. OpenGL Shading Language. Redwood City, CA, Addison-Wesley Longman, 2004.

[11] Daniel Sandin, Todd Margolis, Greg Dawe, Jason Leigh, Thomas DeFanti. The Varrier Autostereographic Display, In Proceedings of SPIE, vol. 4297, San Jose, CA, 2001.

[12] Daniel Sandin, Todd Margolis, Jinghua Ge, Javier Girado, Tom Peterka, and Thomas DeFanti. The Varrier ™ Autostereoscopic Virtual Reality Display, In ACM Transactions on Graphics, Proceedings of ACM SIGGRAPH, 24, no. 3, 2005, p. 894-903, 2005.

[13] Alexander Schmidt and Armin Grasnick. Multi-viewpoint Autostereoscopic Displays from 4D-Vision. In Proceedings of SPIE Photonics West 2002: Electronic Imaging, San Jose, CA, 2002.

[14] David Montgomery, Graham Woodgate, Adrian Jacobs, Jonathan Harrold, and David Ezra. Analysis of a Flat Panel Display System Convertible Between 2D and Autostereoscopic 3D Modes. In Proceedings of SPIE, 4297, p. 148-159, 2001.

[15] Ken Mashitani, Goro Hamagishi, Masahiro Higashino, Takahisa Ando, and Satoshi Takemoto. Step Barrier System Multi-View Glassless 3-D Display. In Proceedings of SPIE-IS&T Electronic Imaging, 5291, p. 265-272, 2004.

[16] Richard Moseley, Graham Woodgate, and David Ezra. Observer Tracking Directional Display. United States Patent 6377295, 2002.

[17] Neil Dodgson, John Moore, Stewart Lang, Graham Martin, and Peter Canepa. A 50" Time-Multiplexed Autostereoscopic Display. In Proceedings of SPIE Electronic Imaging, 3957, 2000.

[18] Ken Perlin, Salvatore Paxia, and Joel Kollin. An Autostereoscopic Display. In Proceedings of ACM SIGGRAPH 2000, Computer Graphics Proceedings, Annual Conference Series, p. 319-326, 2000.

[19] Tom Peterka, Robert Kooima, Javier Girado, Jinghua Ge, Daniel Sandin, Andrew Johnson, Jason Leigh, Jurgen Schulze, and Thomas DeFanti. Dynallax: Solid State Dynamic Parallax Barrier Autostereoscopic VR Display. In Proceedings of IEEE Virtual Reality 2007, p. 155-162, 2007.

[20] Tom Peterka, Robert Kooima, Javier Girado, Jinghua Ge, Daniel Sandin, and Thomas DeFanti. Evolution of the Varrier Autostereoscopic VR Display. In IS&T/SPIE Electronic Imaging 2007 Conference Proceedings, 2007.

[21] NEC Inc.. http://www.necdisplay.com, 2007.

[22] S. Lee, C. Mao, and K. Johnson. Fast-switching liquid-crystal-on-silicon microdisplay with framebuffer pixels and surface-mode optically compensated birefringence. In Optical Engineering 45, no. 12, 2006.

[23] Ken Perlin, C. Poultney, J. Kollin, D. Kristjansson, and S. Paxia. Recent Advances in the NYU Autostereoscopic Display. In Proceedings of SPIE 4297, 2001.

Tom Peterka is a postdoctoral appointee at Argonne National Laboratory. He received his Ph.D. from the University of Illinois at Chicago (UIC) and conducted his graduate research at the Electronic Visualization Laboratory at UIC. His dissertation focused on dynamic barrier autostereoscopy, and he has worked on autostereoscopic display systems since 2003. Currently his work concentrates on interactive ultra-scale massively parallel visualization with autostereo displays serving as an interface between scientists and data.

Robert Kooima is a Ph.D. student at the University of Illinois at Chicago and a graduate research assistant at the Electronic Visualization Laboratory. He received B.S. and M.S. degrees from the University of Iowa in 1997 and 2001. His research interests include real-time 3D graphics and large-scale display technology.

Daniel J. Sandin is an internationally recognized pioneer of electronic art and visualization. He is director emeritus of the Electronic Visualization Laboratory and a professor emeritus in the School of Art and Design at the University of Illinois at Chicago. As an artist, he has exhibited worldwide and has received grants from the Rockefeller Foundation, the Guggenheim Foundation, the National Science Foundation and the National Endowment for the Arts.

Andrew Johnson is an associate professor in the department of computer science and a member of the Electronic Visualization Laboratory at the University of Illinois at Chicago. His research interests focus on interaction and collaboration in advanced visualization environments.

Jason Leigh is an associate professor of computer science and co-director of the Electronic Visualization Laboratory at the University of Illinois at Chicago. His areas of interest include: developing techniques for interactive, remote visualization of massive data sets and for support of long-term collaborative work in amplified collaboration environments. Leigh has led EVL's tele-immersion research agenda since 1995, after developing the first networked CAVE application in 1992. He is co-founder of the GeoWall Consortium and visualization lead in the National Science Foundation's OptIPuter project.

Thomas A. DeFanti, while at the University of Illinois at Chicago, was a director of the Electronic Visualization Laboratory. At the University of California, San Diego, DeFanti is a research scientist at the California Institute for Telecommunications and Information Technology (Calit2). He shares recognition along with EVL director Daniel J. Sandin for conceiving the CAVE virtual reality theater in 1991. Striving for a more than a decade now to connect high-resolution visualization and virtual reality devices over long distances, DeFanti has collaborated with Maxine Brown of UIC to lead state, national, and international teams to build the most advanced production-quality networks available to academics.